arXiv:1404.5757v1 [math.CO] 23 Apr 2014

UNIVERSAL GRAPHS WITH A FORBIDDEN SUBGRAPH: BLOCK PATH SOLIDITY

GREGORY CHERLIN
DEPARTMENT OF MATHEMATICS,
RUTGERS UNIVERSITY, U.S.A.
AND
SAHARON SHELAH
DEPARTMENT OF MATHEMATICS
HEBREW UNIVERSITY, JERUSALEM, ISRAEL
AND RUTGERS UNIVERSITY, PISCATAWAY, NJ 08854

Abstract. Let C be a finite connected graph for which there is a countable universal C-free graph, and whose tree of blocks is a path. Then the blocks of C are complete. This generalizes a result of Füredi and Komjáth, and fits naturally into a set of conjectures regarding the existence of countable C-free graphs, with C an arbitrary finite connected graph.

Introduction

The main theorem. The general problem to be considered here is the following.

Problem 1 (Universality with 1 Forbidden Subgraph)**.** *Let C be a finite connected graph. When is there a universal C-free graph?*

We anticipate that the problem as posed has an explicit solution, but not a very simple one. The present paper is devoted toward the proof of one concrete result which is part of a general plan of attack on the problem. That result reads as follows.

Theorem 1 (Main Theorem)**.** *Let C be a finite connected block path, and suppose there is a (weakly) universal C-free graph. Then the blocks of C are complete.*

The terminology will need to be explained in detail, but we first lay out the context for the result. The general plan of attack referred to consists mainly of the following two conjectures.

Conjecture 1. *Let C be a finite connected graph allowing a countable universal C-free graph. Then*

2010 *Mathematics Subject Classification.* Primary 05C60; Secondary 05C63,03C15.
First author supported by NSF Grant DMS-1101597.
The second author would like to thank the NSF, grant no. DMS 1101597, for support of this paper. Publication 1033.

- *(Solidity Conjecture) The blocks of C are complete;*
- *(Pathlike Conjecture) C may be obtained from a block path by adjoining paths to some of the vertices, with at most one path joined to each vertex.*

Our main theorem may be phrased as follows: the second conjecture implies the first. Given both conjectures, what then remains is to analyze the graphs referred to as "pathlike" here, under the assumption that all blocks are complete.

Recall that the blocks of a graph are its maximal 2-connected subgraphs, and that there is associated to any connected graph a "tree of blocks" whose vertices are the blocks and cut vertices of the original graph, with edges corresponding to incidence. By a *block path* we mean a graph whose tree of blocks forms a path. Such graphs are certainly pathlike, but we also will need to consider block paths with "whiskers," as described in the second conjecture.

Now we review the terminology relating to universality.

Definition (Universality)**.** With C a given *forbidden* subgraph, a graph Γ is *C-free* if Γ contains no subgraph isomorphic to C.

If Γ is countable and C-free, then Γ is

- *strongly universal* if every countable C-free graph is isomorphic to an induced subgraph of Γ;
- *weakly universal* if every countable C-free graph is isomorphic to a subgraph of Γ (such an isomorphism will be called an embedding, or more explicitly, an embedding as a subgraph).

The most natural notion of universality is the strong version. In proving the existence of universal graphs we always aim at strong universality. But when proving nonexistence it is more satisfactory to prove the nonexistence of a *weakly* universal graph, and as the marginal cost of this refinement tends to be low, this is what we usually aim at. In particular, we have stated the main theorem in this sharper form. We also stress that we work with the class of countable C-free graphs throughout. Other cases are of interest but involve different issues and a broad range of techniques (cf. [KP91, Dž05]).

The proof of the main theorem is not very complicated. We use two techniques: the *hypergraph method* of Füredi and Komjáth [FK97], and two *pruning techniques*, one of which has made an appearance in [CS05] in a more specialized form. The pruning techniques allow a certain kind of inductive argument to be carried out, typically reducing a riot of uncontrollable cases to a menagerie of manageable cases. We expect to make substantial further use of the pruning method again in more complicated situations, so one of our goals here is to set this up for future reference.

We view the explicit classification of all finite connected constraint graphs C allowing a universal C-free graph as an ambitious classification project. But what really interests us is the following more qualitative question.

Problem 2 (Universality with finitely many Forbidden Subgraphs)**.** *Is there an algorithm which will decide, for any finite set $\mathcal{C}$ of finitely many connected graphs, whether there is a universal $\mathcal{C}$-free graph?*

This problem can be approached from many sides. It can be shown that if one forbids not just subgraphs, but induced subgraphs, then one arrives at a still more general question, but one for which there is a direct proof of *algorithmic undecidability* by encoding Wang's domino problem [Ch11]. On the other hand if one restricts attention to graphs of bounded degree—that is, one allows a finite set of forbidden connected graphs, one of which is a star—then the problem becomes decidable, though one would not expect a completely explicit classification at that level of generality.

Some prior work. For the purposes of the present article, the outstanding prior result is due to Füredi and Komjáth.

Fact (2-Connected Constraints, [FK97])**.** *Let C be a finite 2-connected graph for which there is a weakly universal C-free graph. Then C is complete.*

Conversely, it is very well known that when C is complete, there is indeed a universal C-free graph [He71].

We will present the Füredi–Komjáth *hypergraph technique* in detail, and make further use of it, in §1. This technique uses a certain hypergraph with good properties as a template for a construction.

Füredi and Komjáth stated a more general result. Call a block of C *distinguished* if it embeds as a subgraph in no other block (in particular it may not be isomorphic to another block). They show that the existence of a C-free universal graph forces the distinguished blocks of C to be complete. As it turns out, we need a different variation on their theme.

A result whose relevance to our current concerns is less obvious is the following.

Fact (Tree constraints, [CT07, CS05])**.** *Let C be a finite tree for which there is a universal C-free graph. Then C is either a path, or may be obtained from a path by adjoining one further edge.*

The point here is the method of proof used, rather than the result obtained. The analysis made use of a form of what we will refer to as *corner pruning* below. We also introduce a second pruning technique we call *symmetric local pruning*, which has not previously seen the light of day.

The pruning method supports an inductive analysis, which in the case of one forbidden tree led us to consider 14 possible cases as the base of the induction. One may wonder whether one can deal with an arbitrary constraint graph by such a method. We believe this is feasible.

A larger view of the wide range of problems connected with universal graphs may be found in the surveys [KP91, Dž05]. From the point of view of those surveys, the restriction to *countable* universal graphs determined by a *finite* set of constraints covers many cases of interest, but by no means all.

The article [Ch11] is in large part a survey of the universality problem restricted to the countable setting, allowing only a finite number of constraints. By restricting to this narrower context we reach a clear algorithmic problem which seems to us to raise a fundamental issue about the subject: namely, is it an art or a science?

1. The Füredi–Komjáth Hypergraph Construction

The proof of the main theorem will be given in three stages, where at each stage we first introduce a general technique, then apply that technique to analyze a hypothetical minimal counterexample to our main theorem.

In the present section we will discuss the Füredi–Komjáth hypergraph construction and apply it to obtain the following special case of our main theorem. Since Füredi and Komjáth alread used their method to treat the case of a single block, this amounts to finishing the base case of an inductive analysis to be taken up further in succeeding sections.

Lemma 1.1. *Let C be a finite connected graph with exactly two blocks (B_1, B_2), and with*

$$|B_1| \leq |B_2|$$

Suppose that there is a weakly universal C-free graph. Then B_1 is a complete graph.

We first present the hypergraph construction of Füredi and Komjáth [FK97].

1.1. **The Hypergraph construction.**

Definition 2.

A *cycle* in a hypergraph is a sequence of distinct vertices and edges $(v_0, E_0, v_1, \ldots, v_{n-1}, E_{n-1})$ with $v_i, v_{i+1} \in E_i$—here we take $(i+1)$ mod n, and $n \geq 2$.

The *girth* of a hypergraph is the length n of the shortest cycle (or ∞).

A hypergraph is *k-uniform* if its hyperedges consist of k points.

Note that if the girth of a hypergraph is greater than 2 then distinct hyperedges meet in at most one vertex. The following is a slight variation on a result of [FK97].

Lemma 1.2. *For any k, g with $k \geq 2$ there is some $N = N(k, g)$ and a k-uniform hypergraph of girth at least g on the vertex set $\mathbb{N}$, with hyperedges E_i ($i \in \mathbb{N}$) satisfying*

$$N + (i-1), N + i \in E_i \subseteq [0, \ldots, N + i]$$

for all i.

Proof. We will impose an additional condition on the hypergraph:

No vertex belongs to more than k hyperedges

We proceed inductively. For the moment, let N be arbitrary, and suppose we have constructed a k-uniform hypergraph on $I_i = [0, \dots, N+i-1]$ with hyperedges $E_0, \dots, E_{i-1}$ satisfying all relevant conditions up to this point. We wish to select E_i.

Let V_i be the set of vertices in I_i belonging to exactly k of the hyperedges E_j for $j < i$. By counting all pairs (u, E_j) with $u \in I_i$, $j < i$, and $u \in E_j$ we find $|V_i| \leq i$. Letting $V_i' = I_i \setminus V_i$ we have $|V_i'| \geq N$. Notice that $N + i - 1 \in V_i'$ since this vertex belongs only to E_{i-1}.

Consider the ordinary graph G_i induced on I_i by taking edges (u, v) whenever $u, v \in E_j$ for some $j < i$. This has vertex degree bounded by $k(k-1)$. Hence in the graph metric we have a bound on the order of balls of radius g, and for N sufficiently large we may select a subset $X \subseteq V_i'$ of cardinality $k-1$, with $N + i - 1 \in X$, and with $d(u, v) > g$ in G_i for $u, v \in X$. Set $E_i = X \cup \{N + i\}$. Then our conditions are all preserved. ∎

1.2. **Application: Lemma 1.1.**

Proof of Lemma 1.1. We have $C = (B_1, B_2)$, where B_1, B_2 are blocks meeting at the unique cut vertex v_* of C. Let $n_i = |B_i|$. We have assumed that

$$n_1 \leq n_2$$

We will suppose that the block B_1 is not complete, and we aim to show that there is no weakly universal C-free graph. As B_1 is not complete we have

$$n_1 \geq 4$$

Let $k = n_1 + 1$, let g be greater than the maximum order of a block of C, and let Γ be a k-hypergraph of girth at least g with the properties of Lemma 1.2. We will label the vertices of Γ as $(u_i : i \in \mathbb{N})$ (so in fact $u_i = i$).

Divide each hyperedge E of Γ into

$$E = E^1 \sqcup E^2 \text{ with } |E^1| = n_1 - 1,\ |E^2| = 2, \text{ and } (\max E) \in E^1$$

Let G_0 denote the graph on Γ in which the induced graphs on each set E^1 are cliques of order $n_1 - 1$, and there are no additional edges. Let G_1 be the graph obtained from G_0 by attaching one clique K_v of order $n_2 + 1$ freely to each vertex v of G_0.

It is clear that G_1 is C-free; we will need a sharper statement given as Claim 2 below.

Claim 1. Let G be a C-free graph, let $n \in \mathbb{N}$, and let $f_1, f_2 : G_1 \hookrightarrow G$ be embeddings of G_1 into G which agree on u_i for $i < n$. Suppose there is a hyperedge E with $u_n = \max E$. Then f_1, f_2 agree on u_n.

We let $u = f_1(u_n)$ and $u' = f_2(u_n)$, which we suppose distinct. Let $B' = f_1[E^1 \setminus \{u_n\}] = f_2[E^1 \setminus \{u_n\}]$ and $B = B' \cup \{u, u'\}$. Then the induced graph on B in G contains all edges except possibly (u, u'). As B_1 is not complete, there is an embedding $j : B_1 \to B$ as a subgraph. Let $v = j(v_*)$. Then $v = f_i(u)$ for some $u \in E^1$ and some $i = 1$ or 2 (or both). Thus

$K = f_i[K_u]$ is a clique of order $n_2 + 1$ containing v and meeting B in at most one other vertex. Let $K' \subseteq K$ be a clique of order n_2 meeting B in $\{v\}$ alone. Then (B, K') contains a copy of C, so C embeds into G, and we have a contradiction. This proves our claim.

Now we extend the graph G_1 in a variety of ways.

Fix an edge e of B_1 containing the cut vertex v_* of C, and let $B_1' = B_1 \backslash e$. Let $\mathcal{E}$ be the set of hyperedges of Γ. For $\varepsilon : \mathcal{E} \to \{0, 1\}$ arbitrary, we will extend G_1 to a graph G_ε as follows.

For each hyperedge E of Γ

(1) If $\varepsilon(E) = 1$ then attach an edge to E^2;
(2) If $\varepsilon(E) = 0$ then attach a copy of B_1' to E^2 with the ends of the deleted edge e corresponding to the vertices of E^2.

Claim 2. Each of the graphs G_ε is C-free.

For E a hyperedge of Γ, let $\hat{E}$ be the portion of G_ε supported by E, namely the induced graph on the union of E together with all the attached cliques K_v ($v \in E$), and also (when $\varepsilon(E) = 0$) the attached copy of B_1'. By the choice of the girth g, if $f : C \to G_\varepsilon$, then the image of either block $f[B_i]$ must lie in one of the sets $\hat{E}$ (not necessarily unique, since the attached cliques K_v are shared by several of the $\hat{E}$).

As neither block B_1, B_2 can be mapped into a copy of B_1', or into a clique of order $n_1 - 1$, or into a single edge (on E^2), these blocks must both go into an attached clique K_v. But then they must go into the same clique K_v, and as $n_1 > 2$ this is impossible.

Now the concluding argument follows a well worn path: namely, we have constructed uncountably many suitably incompatible C-free graphs, and therefore no countable C-free graph can be weakly universal. We give this final argument in detail.

Suppose there is a countable weakly universal C-free graph G, and choose embeddings $f_\varepsilon : G_\varepsilon \to G$ of each G_ε as a subgraph of G. Let N be the parameter associated with the hypergraph Γ, with the property:

$$\text{For any } n \geq N, \text{ there is a hyperedge } E \text{ of } \Gamma \text{ with } \max E = u_n$$

As C is countable, there will be a pair of distinct ε_1, ε_2 for which the corresponding embeddings agree on u_i for $i < N$, and hence agree for all i in view of Claim 1.

Now consider a hyperedge E with $\varepsilon_1(E) \neq \varepsilon_2(E)$. We may suppose

$$\varepsilon_1(E) = 0;\ \varepsilon_2(E) = 1$$

Then the image of E^2 in G contains an edge. Now consider the copy B_1^* of B_1' attached to E^2 in G_{ε_1}, and the vertex v^* of B_1^* corresponding to the cut vertex of C. Write f for f_{ε_1}. Then $f[B_1^* \cup K_{v^*}]$ contains a copy of $C \backslash e$ where e, the deleted edge of B_1, corresponds to the edge on $f[E^2]$ in G. That is, we have now embedded C into G, arriving at a contradiction. ∎

2. Corner Pruning

We now aim at the following reduction of our main theorem.

The *length* of a block path will be defined as the number of blocks.

Lemma 2.1. *Suppose that $C = (B_1, \dots, B_\ell)$ is a block path of length ℓ allowing a weakly universal graph, and having a block which is not complete. Suppose further that the length ℓ is minimal, that $|B_1| \leq |B_\ell|$, and that if B_ℓ is isomorphic to a subgraph of B_1, then B_ℓ is isomorphic to B_1. For $i < \ell$, let v_i be the cut vertex between B_i and B_{i+1}. Then the following hold.*

(1) B_1 is not complete
(2) $\ell \geq 3$
(3) B_i is complete for $1 < i < \ell$
(4) If B_ℓ does not embed in B_1, then B_ℓ is complete
(5) The induced subgraph $(B_1, B_2 \setminus \{v_2\})$ of C embeds into the induced subgraph $(B_2 \setminus \{v_1\}, B_3, \dots, B_\ell)$.

We view the last of these conditions is a weak form of symmetry. One case to keep in mind is that in which the length is 3 and B_1, B_3 are isomorphic.

2.1. The method of corner pruning.

Definition 3.

1. A *segment* of a graph C is a connected subgraph which is a union of blocks.

2. A *corner* C_v of a graph C is a segment of the form $\{v\} \cup C'$ where v is a cut vertex and C' is one of the connected components of $C \setminus \{v\}$. Note that C_v contains a unique block B of C with $v \in B$, and that the pair (v, B) determines the corner. We call v the *root* of C_v, and B its *root block*. Note that a corner will frequently be treated as a graph with base point v (or briefly: a *pointed graph*). For pointed graphs we use the notation

$$(v, C)$$

In particular, we may consider embeddings of one corner into another either as a subgraph, or as a pointed subgraph.

Definition 4 (Pruning)**.** Let Σ be a set of pointed graphs, C a graph, and $\mathcal{C}$ a finite set of graphs.

1. A corner C_v of C is *pruned by* a pointed graph (u, S) if there is an embedding of (v, C) into (u, S) as a pointed subgraph.

2. The Σ-pruned graph C_Σ is the graph obtained from C by deleting the set of vertices in

$$\bigcup_{(u,S)\in\Sigma} \{C_v \setminus \{v\} \mid (v, C_v) \text{ is pruned by } (u, S)\}$$

Thus we do not delete the base point of a pruned corner, only the remainder.

3. $\mathcal{C}_\Sigma = \{C_\Sigma \mid C \in \mathcal{C}\}$

4. Generally we write C' and $\mathcal{C}'$ for the pruned graph or set of graphs, after specifying the set Σ.

We focus here on single constraints C and we prune by a single *minimal* corner. But there is a distinction even in simple cases between pruning by a set of corners taken together, and pruning by a sequence of corners individually and consecutively. We will not require the notion in full generality for our present purposes, but this is likely to come into play in more elaborate analyses.

In [CS05], we dealt with the case of one forbidden tree, and we pruned only leaves. The simplest case of a corner would be a block occurring as a leaf in the tree of blocks. The proof of the next result is much the same as in [CS05].

Lemma 2.2 (Pruning Induction)**.** *Let $\mathcal{C}$ be a set of graphs, and Σ a set of pointed graphs. If there is a countable universal $\mathcal{C}$-free graph (in either the weak or strong sense) then there is a countable universal $\mathcal{C}_\Sigma$-free graph.*

Proof. If G is any graph, define $G^+(\Sigma)$ as the graph obtained by freely adjoining infinitely many copies of each pointed graph (u, S) in Σ to each vertex v of G, identifying u and v. (In particular for v a single vertex, viewing v as a trivial graph, we have the notation $v^+(\Sigma)$.)

If G is $\mathcal{C}_\Sigma$-free, then $G^+(\Sigma)$ is $\mathcal{C}$-free: if $C \in \mathcal{C}$ embeds as a subgraph into $G^+(\Sigma)$, then the part of C lying in G would contain C_Σ.

So now suppose there is a weakly or strongly universal $\mathcal{C}$-free graph Γ, and let Γ_Σ be the induced graph on the set

$$\{v \in \Gamma \mid v^+(\Sigma) \text{ embeds into } \Gamma \text{ over } v\}$$

We will check that Γ_Σ is $\mathcal{C}_\Sigma$-free universal, in the same sense.

Certainly Γ_Σ is $\mathcal{C}_\Sigma$-free, as otherwise we could reattach the pruned corners in Γ.

So let G be $\mathcal{C}_\Sigma$-free, and embed $G^+(\Sigma)$ into Γ, either as a subgraph or as an induced subgraph, as the case may be. Then G goes into Γ_Σ. ∎

Remark 5. There is also some use for a more sensitive notion of pruning, in which we prune only segments which embed into the given pointed graphs as *induced subgraphs.* But this would be relevant only in proving the nonexistence of strongly universal graphs, while we aim at proving nonexistence for weakly universal graphs.

The classification of forbidden trees C allowing a universal C-free graph comes down to the following, by leaf pruning.

Fact ([CS05])**.** *Let T be a tree which becomes either a path or a near path on removal of its leaves. If there is a weakly universal T-free graph, then T is a path or a near path.*

This amounts to the base of an induction, and occupies the bulk of [CS05]; the rest of the induction is purely formal, as we have seen.

We will now proceed similarly with the proof of our main theorem. Corner pruning plus the hypergraph construction will not do everything, but will leave a definite configuration suitable for further analysis by a third method.

It should be noted that the hypergraph construction could do a good deal more than we have done with it—but only via partial overlap with cases handled more thoroughly by pruning.

2.2. **Application: Lemma 2.1.**

Proof of Lemma 2.1. We are supposing that C is a block path with blocks $(B_1, \ldots, B_\ell)$, allowing a countable weakly universal C-free graph, and having some incomplete block, with the length ℓ minimal.

We may suppose further that

$$|B_1| \leq |B_\ell|$$

and that if B_ℓ embeds into B_1, then the two blocks are isomorphic.

By the result of Füredi and Komjáth, $\ell \geq 2$.

The block B_1 is a corner of C and we may prune it. Our assumptions on B_1, B_ℓ imply that this pruning will remove only B_1 and possibly B_ℓ, the latter only if B_1 and B_ℓ are isomorphic. What is left after pruning is a shorter block path with similar properties, so by the minimality of ℓ all of the remaining blocks are complete: that is, B_i is complete for $1 < i < \ell$, and also B_ℓ is complete if B_1 and B_ℓ are not isomorphic.

Since C has some incomplete block, it follows that B_1 is incomplete. So at this point we have

> B_1 is incomplete
> $\ell \geq 3$ (Lemma 1.1)
> B_i is complete for $1 < i < \ell$; and for $i = \ell$ unless $B_1 \cong B_\ell$

So points $(1-4)$ of Lemma 2.1 have been verified.

Our final claim (5) is that we have an embedding of $(B_1, B_2 \setminus \{v_2\})$ into $(B_2 \setminus \{v_1\}, B_3, \ldots, B_\ell)$, where v_i denotes the cut vertex between B_i and B_{i+1}.

It will be useful to bear in mind that the corners of C are its terminal segments

$$R_j = (B_j, B_{j+1}, \ldots, B_\ell) \text{ with base point } v_{j-1}, \text{ and}$$
$$L_j = (B_j, B_{j-1}, \ldots, B_1) \text{ with base point } v_j,$$

where we write "R" and "L" to suggest "right" and "left".

Now we prune the corner $R_3 = (B_3, \ldots, B_\ell)$ with base point v_2. If B_1 remains after pruning, then by the minimality of the length ℓ, B_1 is complete, a contradiction.

If B_1 does not remain after pruning, then it meets, and hence lies within, some corner pruned by R_3, which must be of the form $L_j = (B_j, \ldots, B_1)$ with base point v_j (taking the blocks in reverse order). So L_j embeds into the corner R_3, with v_j corresponding to v_2.

Suppose first that $j = 1$. Then we have an embedding of B_1 into B_3 with v_1 going to v_2. As B_2 is complete, we may extend this to an embedding of $(B_1, B_2 \setminus \{v_2\})$ into $(B_2 \setminus \{v_1\}, B_3)$, taking v_2 to v_1, proving our claim.

Now suppose $j > 1$. Then our embedding takes (B_2, B_1) into $(B_3, \ldots, B_\ell)$ with v_2 fixed if $j = 2$, and with v_2 not in the image if $j > 2$. So claim (5) holds in either case. ■

We will see in the next section that the weak symmetry condition (5) allows another kind of pruning.

3. Local and symmetric Pruning

In this section we aim to complete the proof of the main theorem by dealing with the configuration described in Lemma 2.1. We introduce another, more subtle, pruning technique. At this point we will confine our theoretical discussion to the case of a single constraint, though no doubt this tool is useful in greater generality.

3.1. **The method of local pruning.**

Local pruning is a way of removing a single corner. We first give the definition in pragmatic terms, and then look for reasonable conditions sufficient for its application.

Definition 6. Let (v, C_v) be a corner of the finite connected graph C. Let C_v^+ be the union of the other corners of C rooted at v, and $C_v^- = C_v \setminus \{v\} = C \setminus C_v^+$.

1. For any graph H, the graph $\hat{H} = H *_v C_v^+$ is the suspension of H with an attached copy of C_v^+, constructed as follows.

(a) Take the disjoint union $H \sqcup C_v^+$

(b) Connect v to every vertex of H by an edge.

2. (v, C_v^+) is said to be *detachable* if the following holds.

Whenever H is C_v^--free, then $\hat{H}$ is C-free.

Lemma 3.1. *Let C be a finite connected graph, and (v, C_v) a corner with (v, C_v^+) detachable. If there is a countable universal C-free graph, in either the weak or strong sense, then there is a countable universal C_v^--free graph, in the same sense.*

Proof. Let Γ be a universal C-free graph, in one of the two senses.

For each embedding h of (v, C_v^+) into Γ as a subgraph, let Γ_h be

$$\{u \in \Gamma \mid u \notin h[C_v^+] \text{ and } (u, h(v)) \text{ is an edge}\}$$

Let $\Gamma_0 = \bigsqcup_h \Gamma_h$ (a *disjoint union*) with h varying over weak embeddings (as subgraphs) if we are in the weak case, or over strong embeddings (as induced subgraphs) in the strong case. We claim that Γ_0 is countable universal C_v^--free, in the corresponding sense.

As C_v^- is connected and the individual Γ_h are C_v^--free, the graph Γ_0 is C_v^--free. Now we check the universality.

If H is any countable C^--free graph then we form the extension $\hat{H} = H *_v C_v^+$ and by hypothesis $\hat{H}$ is C-free, hence embeds into Γ. Then this embedding takes H into the corresponding induced subgraph Γ_h in Γ_0. ■

3.2. **A special case: Symmetric Local Pruning.** Now we become more concrete, in the context of block paths. We continue to work with the notation of the previous section.

In analyzing detachability, we must pay particular attention to "improbable" embeddings of a given graph C in some graph of the form $\hat{H}$.

Lemma 3.2 (Symmetric Local Pruning)**.** *Let C be a block path, B a block of C containing two cut vertices u, v, and let L_u, R_u, L_v, R_v be the corners rooted at u and v respectively, with R_u and L_v the ones containing the block B. Suppose that $L_v \setminus \{v\}$ embeds into $R_u \setminus \{u\}$. Then (v, R_v) is detachable.*

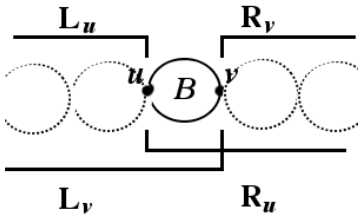

Remark 7. Here the block B is common to L_v and R_u, so one possible type of embedding would involve a symmetry over B. In practice other types of embedding will also occur, so the concept of "symmetry" used here is very general.

In our notation, L and R stand for *left* and *right*. It is helpful to think of the tree of blocks, which is a path, ordered so that $u < B < v$. Note that $R_v = L_v^+$. In particular detaching R_v leaves $L_v^- = L_v \setminus \{v\}$.

Proof of Lemma 3.2. We suppose that H is an L_v^--free graph, and we set

$$\hat{H} = H *_v (v, R_v)$$

Our claim is that $\hat{H}$ is C-free.

We will write (v_0, R_0) for the copy of (v, R_v) in $\hat{H}$, and fix an isomorphism

$$\iota : (v_0, R_0) \cong (v, R_v)$$

Suppose toward a contradiction that

$$f : C \cong \hat{C} \subseteq \hat{H}$$

where the inclusion is as a subgraph.

If $f[R_v] \subseteq R_0$ then $f[R_v] = R_0$ and $f[L_v^-] \subseteq H$, a contradiction. So

$$f[R_v] \not\subseteq R_0$$

Case 1. Suppose first that $v_0 \notin f[R_v]$.

In this case we must have $f[R_v]$ contained either in $R_0 \setminus \{v_0\}$ or in H, and the first alternative is out of the question. Thus

$$f[R_v] \subseteq H$$

In particular, $f(v) \in H$ and $f[B] \subseteq H \cup \{v_0\}$.

If $f[B] \subseteq H$ then $f[R_u] \subseteq H$ and L_v^- embeds into H, a contradiction.

So suppose $v_0 \in f[B]$ and $v_1 = f^{-1}(v_0)$. Then as $v_1 \neq v$ and B is complete, we have $(R_u \setminus \{u\}) \cong (R_u \setminus \{v_1\})$. But $f[(R_u \setminus \{v_1\})] \subseteq H$ and thus L_v^- embeds into H, a contradiction.

Case 2. Suppose that $v_0 = f(v)$.

Then $f[L_v^-]$ is contained either in H or in $R_0 \setminus \{v_0\}$. As the first alternative is excluded by hypothesis, we have

$$f[L_v^-] \subseteq R_0 \setminus \{v_0\}$$

Since $f[R_v] \not\subseteq R_0$, we have

$$f[R_v] \subseteq H \cup \{v_0\}, \text{ with } f(v) = v_0$$

So $f\iota f[L_v^-] \subseteq H$, a contradiction.

Case 3. Suppose $v_0 = f(v_1)$ with $v_1 \in R_v \setminus \{v\}$.

This is the most delicate case.

As $v \notin f[L_v^-]$ we again have $f[L_v^-]$ contained in H or $R_0 \setminus \{v_0\}$, with the first alternative ruled out by hypothesis. So we have

$$f[L_v^-] \subseteq R_0 \setminus \{v_0\}$$

In particular

$$f(v) \in R_0 \setminus \{v_0\}$$

In what follows we are mainly concerned about the relation of $\iota f(v)$ to v_1.

Let S be the smallest segment of R_v containing v and $\iota f(v)$.

Suppose first that

$$f[S] = \iota^{-1}[S] \tag{1}$$

Then $f[S] \subseteq R_0$, and $v_0 \in f[S]$, so since $f[R_v] \not\subseteq R_0$ we find $f[R_v \setminus S] \subseteq H$.

Now $f[L_v] \cap \iota^{-1}[S] = f[L_v \cap S] = \{f(v)\}$, so $\iota f[L_v] \cap S = \{\iota f(v)\}$. Thus $\iota f[L_v^-] \subseteq R_v \setminus S$, and so $f\iota f[L_v^-] \subseteq f[R_v \setminus S] \subseteq H$, and we have an embedding of L_v^- into H, for a contradiction.

There remains the alternative

$$f[S] \neq \iota^{-1}[S] \tag{2}$$

and hence

$$f[S] \not\subseteq \iota^{-1}[S]$$

Next we claim

$$f(\iota f(v)) \in H$$

Otherwise, we have $f[S] \subseteq R_0$. But $\iota f(v) \in S \cap \iota f[S]$, $\iota f[S] \not\subseteq S$, so $\iota f[S]$ contains the cut vertex between S and $R_v \setminus S$ and the adjacent block of $R_v \setminus S$. This then forces $f[R_v] \subseteq R_0$, a contradiction. So $\iota f(v) \in H$.

In particular, v_0 is a cut vertex of $f[S]$, and thus

$$v_1 \text{ is a cut vertex of } S.$$

Let S_1 be the segment from v to v_1 in R_v.

Now $f[R_v \setminus S] \subseteq H$, and $f(v_1) = v$, so

$$f[R_v \setminus S_1] \subseteq H$$

Now $\iota^{-1}(v_1)$ is a cut vertex of R_0 lying between v_0 and $f(v)$. Hence $\iota^{-1}(v_1) \in f[S]$, and $v_1 \in \iota f[S]$. As $v_1 \neq v$, we have

$$v_1 \notin \iota f[L_v^-]$$

But $\iota f(v) \in \iota f[L_v]$, so $\iota f[L_v^-] \subseteq R_v \setminus S_1$. So $f\iota f[L_v^-] \subseteq H$ and again we have a contradiction. ■

3.3. **Application: The Main Theorem.**

Proof of Theorem 1. We suppose toward a contradiction that

$$C = (B_1, \ldots, B_\ell)$$

is a block path with at least one incomplete block, allowing a weakly universal C-free graph, and with the length ℓ minimized.

We may suppose

$$|B_1| \leq |B_\ell|$$

We claim that we may also suppose that one of the following two conditions applies

- $B_1 \cong B_\ell$
- B_ℓ is not isomorphic to a subgraph of B_1

If $|B_1| < |B_\ell|$ this is clear, while if $|B_1| = |B_\ell|$, we are free to switch the roles of B_1 and B_ℓ. So unless B_1 and B_ℓ are isomorphic, we may suppose that B_ℓ does not embed isomorphically into B_1.

So we arrive at the conditions of Lemma 2.1, and in particular at the conclusions that $\ell \geq 3$, B_1 is not complete, and that $L_2^- = (B_1, B_2 \setminus \{v_2\})$ embeds into $(B_2 \setminus \{v_1\}, \ldots, B_\ell)$. Taking $B = B_2$ in Lemma 3.2, we find that $R_3 = (B_3, \ldots, B_\ell)$ is detachable. Thus there is a weakly universal L_2^--free graph. So by the case $\ell = 2$ (or $\ell = 1$ if $|B_2| = 2$), the block B_1 is complete. This contradiction completes the proof. ■

References

[Ch11] G. Cherlin Two problems on homogeneous structures, revisited, *in* Proceedings of the AMS-ASL Special Session on Model Theoretic Methods in Finite Combinatorics, January 5-8, 2009, in Washington, DC.; pp. 319-416 in Contemporary Mathematics, AMS, Providence, RI, 2011.

[CS05] G. Cherlin and S. Shelah, Universal graphs with a forbidden subtree. *J. Comb. Theory, Series B* **97** (2007), 293–333.

[CT07] G. Cherlin and L. Tallgren, Universal graphs with a forbidden near-path or 2-bouquet, *J. Graph Theory* **56** (2007), 41–63.

[Dž05] M. Džamonja, Club guessing and the universal models. *Notre Dame J. Formal Logic* **46** (2005), 283–300.

[FK97] Z. Fűredi and P. Komjáth, On the existence of countable universal graphs. *J. Graph Theory*, **25** (1997), 53–58.

[He71] C. Ward Henson, A family of countable homogeneous graphs. *Pacific J. Math.* 38:69–83, 1971.

[KP91] P. Komjáth and J. Pach, Universal elements and the complexity of certain classes of infinite graphs. *Discrete Math.* **95** (1991), 255–270.